arXiv:1610.08974v1 [stat.AP] 27 Oct 2016

Phylogenetic Dirichlet-multinomial model for microbiome data

Yunfan Tang*, Li Ma†, and Dan L. Nicolae*

*University of Chicago** *and Duke University*†

Abstract

In this paper we introduce the phylogenetic Dirichlet-multinomial (PhyloDM) model for investigating cross-group differences in microbiome compositions. Traditional Dirichlet-Multinomial (DM) models ignore species relatedness, leading to loss in efficiency and to results that are difficult to interpret. PhyloDM solves these issues by replacing the global model with a cascade of independent local DMs on the internal nodes of the phylogenetic tree. Each of the local DMs captures the count distributions of a certain number of operational taxonomic units (OTU) at a given resolution. Since distributional differences tend to occur in clusters along evolutionary lineages, we design a scan statistic over the phylogenetic tree to allow nodes to borrow signal strength from their parents and children. We also derive a formula to bound the tail probability of the scan statistic, and verify its accuracy through simulations. The PhyloDM model is applied to the American Gut dataset to identify taxa associated with diet habits. Empirical studies performed on this dataset show that PhyloDM achieves a significantly better fit, and has higher testing power than DM.

1 Introduction

Microbiome refers to the full collection of genes of all microbes in a community; for example, all bacteria in a sample from the gut of a healthy individual. The advent of next generation sequencing technologies, such as 454 pyrosequencing and Illumina Solexa, has allowed researchers to investigate the microbiome communities at an unprecedented level of quantification. The focus of this paper is on targeted amplicon sequencing and not on metagenome, but the ideas introduced here can be easily extended to metagenome data. A typical analysis pipeline involves sequencing one or a few of the variable regions of 16s ribosomal RNA, clustering the sequences into operational taxonomic units (OTU), assigning taxonomy to OTUs according to a reference database, and constructing a phylogenetic tree [e.g. Caporaso et al. (2010)]. There has been burgeoning efforts devoted to the study of human microbiome in the past decade, many of which aim at establishing evidence between microbiome and treatment effects or environmental covariates. Examples of such include associating gut microbiome with diet [David et al. (2014)], autism

spectrum disorder [McDonald et al. (2015b)] and hormones [Neuman et al. (2015)]. Besides, a number of these studies are very large-scale initiatives such as Human Microbiome Project [Human Microbiome Project Consortium (2012)] and American Gut [McDonald et al. (2015a)], both of which seek to provide a much broader understanding of the microbial variability. These studies jointly point to the fact that microbiome plays an integral part to our health, and much still remains to be explored in this area.

The vast improvement in experimental methodology contrasts with the much slower development of parametric statistical models to analyze the microbiome data. Typically, the majority of taxa can only be observed in a very small subset of samples, which causes the data table to be highly sparse. In addition, the within-group heterogeneity among samples leads to pronounced overdispersion in taxa proportions. Since standard multinomial distributions fail at capturing these features, researchers have been using Dirichlet-multinomial (DM) as an extension. DM was originally proposed by Mosimann (1962) and introduced into the microbiome context by La Rosa et al. (2012) and Holmes et al. (2012). Although DM is capable of capturing both within-group and cross-group variations across all taxa simultaneously, applying DM to test cross-group variations suffers from a number of drawbacks. First and foremost, DM comparison is conducted on the global level and therefore unable to localize any signal to a subgroup of taxa. Such limitation dramatically inhibits the potential of a vast number of microbiome association studies as most biological explanations of the community difference require pinpointing the taxa that contribute to the variation of interest. Second, when the number of taxa is large, DM test has very little power due to large degrees of freedom. This forces researchers to pool the count data at a much lower resolution, leading to loss of taxonomic information. Third, asymptotic results of method-of-moment or likelihood-ratio-tests fail when a number of taxa has low counts across all samples.

Our remedies to the above concerns is mainly motivated by the Polya Tree process [Lavine (1992)], which characterizes the sample space through a cascade of nested dyadic partitions. Polya Tree has recently been adopted to nonparametrically test the equality of two-sample distributions by Ma and Wong (2011), Chen and Hanson (2014), Holmes et al. (2015) and Soriano and Ma (2016). It turns out that the Polya tree methods are particularly well-suited for microbiome data when the sample is decomposed on a phylogenetic tree. The reason behind this claim is that phylogenetic trees in general represent functional similarities among the OTUs so for each internal node, its descendants will share a certain degree of affinity. Using such binning strategy allows one to assign a meaningful interpretation to the distribution on each internal node. The global two-sample test is therefore represented by a number of independent and biologically relevant constituents. Following this motivation, we propose the phylogenetic Dirichlet-multinomial (PhyloDM) model that applies independent DMs locally on each internal node. In particular, we show that the global DM distribution as used by La Rosa et al. (2012) and Chen and Li (2013) on all taxa counts also admits a sequential decomposition into a set of local DMs, thus belonging to the PhyloDM family. The PhyloDM generalizes the global DM through removing the dispersion constraints on the local distributions. Each of the local DM targets the abundance of the descendants of a particular node and consequently enjoy much

lower degrees of freedom. This breaking down of the global distribution on all taxa counts allows us to test each branch of the phylogeny individually, hence locating the signals to a certain taxonomic rank. The issue of sparse count is also alleviated on higher levels of the phylogeny where counts from large numbers of OTUs are pooled together. The researcher can selectively decide whether to push the model to the next level, depending on the number of counts.

Although standard multiple testing procedures could be applied to results from testing all nodes, we note that it is usually not the best practice to treat each hypothesis as a segregated entity. Soriano and Ma (2016) pointed out that cross-group distributional variations tend to cluster, which causes hypotheses defined on nearby and/or nested windows more likely to be jointly true or false. This observation also holds in the microbiome data; cross-group differences in a certain ancestor node are more frequently accompanied with similar differences in its descendants. To take advantage of this structure and optimize the test power, we adopt ideas from scan tests through constructing a collection of triplet statistics, each incorporating evidence from an internal node in the phylogenetic tree along with its parent and a child node. The maximum of all these triplet statistics is used to test the global null hypothesis. Since the exact distribution of maximum statistic is intractable, we derive an upper and lower bound on its tail probability based on existing results on union probability [e.g., Hunter (1976), Efron (1997) and Taylor et al. (2007)]. Our improved strategy first finds a subset consisting of independent components from the union, followed by bounding the probability of remaining components conditioned on the complement of that subset. A decay rate of the relative error of our approximation is also provided.

The rest of the paper is organized as follows. Section 2 briefly reviews the DM model. Section 3 formulates the PhyloDM model and establishes its relation to the DM. Section 4 develops p-value approximation on the scan statistic for the PhyloDM and verifies the result through simulation. Section 5 applies the PhyloDM model on the American Gut dataset to test the association of gut microbiome with a number of dietary habits. It also empirically demonstrates improvement of PhyloDM over DM through likelihood ratio tests and comparing simulated test power. Section 6 concludes with further discussions on potential PhyloDM extensions.

2 Dirichlet-multinomial for microbiome data

In this section we briefly recap the cross-group testing procedures on microbiome data using the Dirichlet-multinomial model, as presented in La Rosa et al. (2012).

Consider a microbiome dataset with n samples and let Ω be the collection of a total of $K = |\Omega|$ OTUs. Without loss of generality, we assume $\Omega = \{1, 2, ..., K\}$. Each sample is a K-dimensional count vector representing the number of sequences in each of the K OTUs. Let $\boldsymbol{x}_i = (x_{i1}, x_{i2}, ..., x_{iK})$ be the taxa count vector of the ith sample for $i = 1, 2, ..., n$. In addition, define $N_{i\cdot} = \sum_{j=1}^{K} x_{ij}$ to be the total number of sequences in the ith sample, $N_{\cdot j} = \sum_{i=1}^{n} x_{ij}$ to be the total number of sequences in the jth OTU, and

$N_{\cdot\cdot} = \sum_{i=1}^{n} N_{i\cdot} = \sum_{j=1}^{K} N_{\cdot j}$. The Dirichlet-multinomial (DM) model assumes

$$\boldsymbol{q}_i \overset{i.i.d.}{\sim} \text{Dir}(\nu\boldsymbol{\pi}), \quad \boldsymbol{x}_i|\boldsymbol{q}_i \sim \text{Multinomial}(N_{i\cdot}, \boldsymbol{q}_i),$$

where $\boldsymbol{\pi} = (\pi_1, \pi_2, \ldots, \pi_K)$ satisfying $\sum_{j=1}^{K} \pi_j = 1$ and $\pi_j > 0$ denotes the mean taxa proportions and $\nu > 0$ is a dispersion parameter that controls the level of variation across samples. Alternatively one may use $\theta = \frac{1}{1+\nu}$ to parametrize the dispersion so that $0 \leq \theta < 1$. Integrating out the $\boldsymbol{q}_i$ gives

$$f(\boldsymbol{x}_i) = \binom{N_{i\cdot}}{\boldsymbol{x}_i} \frac{\Gamma(\nu)}{\Gamma(N_{i\cdot} + \nu)} \prod_{j=1}^{K} \frac{\Gamma(x_{ij} + \nu\pi_j)}{\Gamma(\nu\pi_j)}. \tag{1}$$

Throughout this paper we use $f(\cdot)$ exclusively to denote the DM probability mass function. When $\nu = \infty$ ($\theta = 0$), the DM degenerates to the standard multinomial distribution. Smaller values of ν indicates larger degrees of overdispersion. Assuming $\boldsymbol{x}_i$'s are independent, the likelihood function is simply the product of probabilities over all samples:

$$\mathcal{L}(\boldsymbol{\pi}, \nu) = \prod_{i=1}^{n} \left[\binom{N_{i\cdot}}{\boldsymbol{x}_i} \frac{\Gamma(\nu)}{\Gamma(N_{i\cdot} + \nu)} \prod_{j=1}^{K} \frac{\Gamma(x_{ij} + \nu\pi_j)}{\Gamma(\nu\pi_j)} \right] \tag{2}$$

As is shown in Weir and Hill (2002), the method of moments (MoM) estimates of the mean proportion $\boldsymbol{\pi}$ and dispersion θ are respectively

$$\hat{\boldsymbol{\pi}} = (\hat{\pi}_1, \hat{\pi}_2, ..., \hat{\pi}_K) \text{ with } \hat{\pi}_j = N_{\cdot j}/N_{\cdot\cdot}$$

$$\hat{\theta} = \frac{\sum_{j=1}^{K} (S_j - G_j)}{\sum_{j=1}^{K} \left(S_j + (N_c - 1)G_j\right)},$$

where $N_c = \frac{N_{\cdot\cdot} - (N_{\cdot\cdot})^{-1} \sum_{i=1}^{n} N_{i\cdot}^2}{n-1}$, $S_j = \frac{\sum_{i=1}^{n} N_{i\cdot}(\hat{\pi}_{ij} - \hat{\pi}_j)^2}{n-1}$, and $G_j = \frac{\sum_{i=1}^{n} N_{i\cdot}\hat{\pi}_{ij}(1-\hat{\pi}_{ij})}{\sum_{i=1}^{n}(N_{i\cdot}-1)}$ with $\hat{\pi}_{ij} = x_{ij}/N_{i\cdot}$.

For hypothesis testing, suppose we collect G groups of observations and the gth group data is given by $\boldsymbol{x}_1^{(g)}, \boldsymbol{x}_2^{(g)}, ..., \boldsymbol{x}_{n_g}^{(g)}$ with $N_{i\cdot}^{(g)} = \sum_{j=1}^{K} x_{ij}^{(g)}$ and $N_{\cdot\cdot}^{(g)} = \sum_{i=1}^{n} N_{i,}^{(g)}$. Similarly we define the gth group parameters as $\boldsymbol{\pi}^{(g)}, \nu^{(g)}$ with $\theta^{(g)} = \frac{1}{1+\nu^{(g)}}$. We wish to test the equality of mean proportion across all groups:

$$H_0 : \boldsymbol{\pi}^{(1)} = \boldsymbol{\pi}^{(2)} = ... = \boldsymbol{\pi}^{(G)} \text{ vs } H_a : \text{otherwise}$$

Let $\hat{\boldsymbol{\pi}}^{(g)}$ and $\hat{\theta}^{(g)}$ be the MoM estimates of $\boldsymbol{\pi}^{(g)}$ and $\theta^{(g)}$, respectively. The cross-group pooled estimate of $\boldsymbol{\pi}$ is $\hat{\boldsymbol{\pi}}^{(Pool)} = \sum_{g=1}^{G} \bar{s}_g \hat{\boldsymbol{\pi}}^{(g)}$ with the following quantity:

$$\bar{s}_g = \frac{(N_{\cdot\cdot}^{(g)})^2 C(\hat{\theta}^{(g)}, N_{\cdot\cdot}^{(g)})^{-1}}{\sum_{r=1}^{G} (N_{\cdot\cdot}^{(r)})^2 C(\hat{\theta}^{(r)}, N_{\cdot\cdot}^{(r)})^{-1}},$$

where

$$C(\hat{\theta}^{(g)}, N_{\cdot\cdot}^{(g)}) = \hat{\theta}^{(g)}\Big(\sum_{i=1}^{n_g}(N_{i\cdot}^{(g)})^2 - N_{\cdot\cdot}^{(g)}\Big) + N_{\cdot\cdot}^{(g)}.$$

Finally, the test statistic is defined as

$$T = \sum_{g=1}^{G}(\hat{\boldsymbol{\pi}}^{(g)} - \hat{\boldsymbol{\pi}}^{(Pool)})^T(\bar{S}_g)^{-1}(\hat{\boldsymbol{\pi}}^{(g)} - \hat{\boldsymbol{\pi}}^{(Pool)}), \tag{3}$$

where $\bar{S}_g$ is a diagonal matrix given by

$$\bar{S}_g = \Big((N_{\cdot\cdot}^{(g)})^2 C(\hat{\theta}_g, N_{\cdot\cdot}^{(g)})^{-1}\Big)^{-1} D(\hat{\boldsymbol{\pi}}^{(Pool)}),$$

and $D(\hat{\boldsymbol{\pi}}^{(Pool)})$ is a diagonal matrix with its diagonal elements given by $\hat{\boldsymbol{\pi}}^{(Pool)}$. The asymptotic distribution of T under H_0 is $\chi^2_{(K-1)(G-1)}$ as $n_g \to \infty$ for all g.

Since the number of OTUs is very large (typically over a thousand), researchers have been grouping OTUs into a certain taxonomic level such as genus [Holmes et al. (2012), Chen and Li (2013)] and class [La Rosa et al. (2012)] in order to effectively reduce the degrees of freedom. However, the additional step of grouping relies on a quite arbitrary choice of taxonomic level. As will be shown in Section 5.1, DM method applied to different taxonomic levels sometimes gives contradictory significance results.

3 Extension of Dirichlet-multinomial on a tree

Aside from its large degrees of freedom, DM does not take species relatedness information into consideration. Such information is often represented by a spanning tree over all categories, which is built from hierarchical clustering according to functional similarities. For microbiome data, phylogenetic tree is a natural candidate to represent species relatedness. In this section we present a novel DM generalization, called phylogenetic Dirichlet-multinomial (PhyloDM), that incorporates functional similarities by modeling the data generative process directly on the phylogenetic tree. We show that this PhyloDM generalization allows us to separately test cross-group differences in each internal node, locating the source of overall difference within particular subgroups of OTUs. Each of the local test, by design, has the benefit of much lower degrees of freedom than an overall DM test.

3.1 Model formulation

Let $\mathcal{T} = (\Omega, \mathcal{I})$ be a rooted phylogenetic tree where the set of OTUs Ω are placed on the leaves and $\mathcal{I}$ is the set of all internal nodes. We represent the elements in $\mathcal{I}$ to be subsets of Ω since each internal node is uniquely identified by the subset of all OTUs underneath it, and vice versa. Each subset of OTU that corresponds to an internal node shares a

hypothetical ancestor along the lineage. Additionally, each leaf node is uniquely identified by a singleton set consisting of that particular OTU.

Figure 1 shows an example of a simple phylogenetic tree over 5 OTUs and 4 internal nodes. This tree has $\Omega = \{1, 2, 3, 4, 5\}$ and $\mathcal{I} = \big\{\{1, 2, 3, 4, 5\}, \{1, 2, 3\}, \{4, 5\}, \{2, 3\}\big\}$.

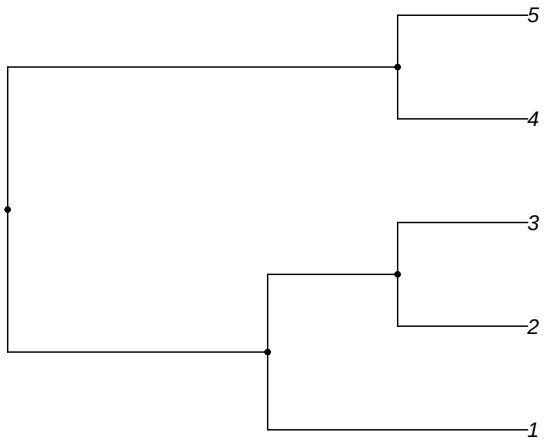

Figure 1: An example of a phylogenetic tree.

Now for $\forall A \in \mathcal{I}$, let $\mathcal{C}(A)$ be the collection of A's children nodes in $\mathcal{T}$. The elements of $\mathcal{C}(A)$ are also subsets of Ω. Also $\forall A \in \mathcal{I} \cup \big\{\{\omega\}|\omega \in \Omega\big\}, A \neq \Omega$, let $R(A)$ denote the parent node of A. In Figure 1, for example, $\mathcal{C}(\{1, 2, 3\}) = \big\{\{1\}, \{2, 3\}\big\}$ and $R(\{1, 2, 3\}) = \{1, 2, 3, 4, 5\} = \Omega$. Notice that certain $\mathcal{C}(A)$'s contain singletons of Ω since some children nodes are leaves. Let $k(A) = |\mathcal{C}(A)|$ be the number of children under A and write $\mathcal{C}(A) = \{\mathcal{C}(A)_1, \mathcal{C}(A)_2, ..., \mathcal{C}(A)_{k(A)}\}$. For each $i = 1, 2, ..., n$ and $j = 1, 2, ..., k(A)$, let

$$x_{ij}(A) = \sum_{\omega \in \mathcal{C}(A)_j} x_{i\omega}$$

be the count of the jth child of A in the ith sample. The count vector associated with A is therefore

$$\boldsymbol{x}_i(A) = \big(x_{i1}(A), x_{i2}(A), ..., x_{ik(A)}(A)\big)$$

with the sum $N_i(A) = \sum_{j=1}^{k(A)} x_{ij}(A) = \sum_{\omega \in A} x_{i\omega}$. It is straightforward to see that $N_i\big(\mathcal{C}(A)_j\big) = x_{ij}(A)$ for $j = 1, 2, ..., k(A)$. In addition, we always have $N_i(\Omega) = N_{i\cdot}$.

Our proposed phylogenetic Dirichlet-multinomial (PhyloDM) separately models the count vector $\boldsymbol{x}_i(A)$ conditional on $N_i(A)$ for each A. Specifically for $\forall A \in \mathcal{I}$,

$$\boldsymbol{q}_{A,i} \overset{i.i.d.}{\sim} \text{Dir}(\nu_A \boldsymbol{\pi_A}), \quad \boldsymbol{x}_i(A)|N_i(A), \boldsymbol{q}_{A,i} \sim \text{Multinomial}(N_i(A), \boldsymbol{q}_{A,i}) \tag{4}$$

where $\nu_A > 0$ is the overdispersion parameter of the counts of A's children and $\boldsymbol{\pi}_A = (\pi_{A,1}, \pi_{A,2}, ..., \pi_{A,k(A)})$ satisfying $\sum_{i=1}^{k(A)} \pi_{A,i} = 1$ denotes their mean proportion. We require

the Dirichlet prior distribution of all A's to be mutually independent. Integrating out $q_{A,i}$ gives

$$f\big(\boldsymbol{x}_i(A)|N_i(A)\big) = \binom{N_i(A)}{\boldsymbol{x}_i(A)} \frac{\Gamma(\nu_A)}{\Gamma(N_i(A)+\nu_A)} \prod_{j=1}^{k(A)} \frac{\Gamma(x_{ij}(A)+\nu_A\pi_{A,j})}{\Gamma(\nu_A\pi_{A,j})}, \tag{5}$$

which ultimately yields

$$f_T(\boldsymbol{x}_i) = \prod_{A\in\mathcal{I}} f\big(\boldsymbol{x}_i(A)|N_i(A)\big) \tag{6}$$

and likelihood function

$$\mathcal{L}_T(\{(\nu_A,\boldsymbol{\pi}_A) : A\in\mathcal{I}\}) = \prod_{i=1}^{n}\prod_{A\in\mathcal{I}} f\big(\boldsymbol{x}_i(A)|N_i(A)\big) \tag{7}$$

Throughout this paper $f_T(\cdot)$ and $\mathcal{L}_T(\cdot)$ are used to denote the PhyloDM probability mass function and likelihood function respectively. The representations in (5) and (6) naturally leads to a top-down generative scheme of the count data on the nodes. We successively apply (5) for each internal node in the order of ascending distance from the root (assuming the root has distance 0 from itself), as each layer of DM models a subset of OTU counts at increased level of resolution conditioned on their sum.

3.2 PhyloDM as a generalization of DM

Interestingly, the DM formula in (1) can be decomposed into the product of $|\mathcal{I}|$ DM probabilities in the form of (6), so our PhyloDM is a strict generalization of DM model on the OTU counts. This is because both the Dirichlet prior and the multinomial probabilities can be factorized over $\mathcal{I}$, i.e.

$$\boldsymbol{q}_i \sim \mathrm{Dir}(\nu\boldsymbol{\pi}) \Leftrightarrow \forall A\in\mathcal{I},\quad \frac{\boldsymbol{q}_i(A)}{\sum_{\omega\in A} q_{i\omega}} \sim \mathrm{Dir}\Big(\nu\sum_{\omega\in A}\pi_\omega\cdot\frac{\boldsymbol{\pi}(A)}{\sum_{\omega\in A}\pi_\omega}\Big) \text{ independently}$$

$$\boldsymbol{x}_i|\boldsymbol{q}_i \sim \mathrm{Multinomial}(N_{i\cdot},\boldsymbol{q}_i) \Leftrightarrow$$

$$\forall A\in\mathcal{I},\quad \boldsymbol{x}_i(A)|\boldsymbol{q}_i,N_i(A) \sim \mathrm{Multinomial}(N_i(A),\frac{\boldsymbol{q}_i(A)}{\sum_{\omega\in A} q_{i\omega}})$$

where we similarly defined

$$q_{ij}(A) = \sum_{\omega\in\mathcal{C}(A)_j} q_{i\omega}, \boldsymbol{q}_i(A) = (q_{i1}(A), q_{i2}(A), ..., q_{ik(A)}(A))$$

$$\pi_j(A) = \sum_{\omega\in\mathcal{C}(A)_j} \pi_\omega, \boldsymbol{\pi}(A) = (\pi_1(A), \pi_2(A), ..., \pi_{k(A)}(A))$$

Alternatively, one can also match the probability mass function on both representations (for which we omit the details):

$$\begin{aligned} f(\boldsymbol{x}_i) &= \binom{N_{i\cdot}}{\boldsymbol{x}_i} \frac{\Gamma(\nu)}{\Gamma(N_{i\cdot}+\nu)} \prod_{j=1}^{K} \frac{\Gamma(x_{ij}+\nu\pi_j)}{\Gamma(\nu\pi_j)} \\ &= \prod_{A\in\mathcal{I}} \binom{N_i(A)}{\boldsymbol{x}_i(A)} \frac{\Gamma(\nu \sum\limits_{\omega\in A}\pi_\omega)}{\Gamma(N_i(A)+\nu \sum\limits_{\omega\in A}\pi_\omega)} \prod_{j=1}^{k(A)} \frac{\Gamma(x_{ij}(A)+\nu\pi_j(A))}{\Gamma(\nu\pi_j(A))} \\ &= \prod_{A\in\mathcal{I}} f(\boldsymbol{x}_i(A)|N_i(A)) \end{aligned} \tag{8}$$

In the cascaded representation of DM, the overdispersion and mean proportion of the counts of A's children are respectively $\nu_A = \nu \sum_{\omega\in A}\pi_\omega$ and $\boldsymbol{\pi}_A = \boldsymbol{\pi}(A)/\sum_{\omega\in A}\pi_\omega$. It is not hard to notice that there is a bijective correspondence between $\boldsymbol{\pi}$ and $\{\boldsymbol{\pi}_A = \boldsymbol{\pi}(A)/\sum_{\omega\in A}\pi_\omega : A \in \mathcal{I}\}$. In addition, all of these local dispersions are governed by a single global ν, which is highly restrictive as it does not allow any node-specific characterization of within-group variation. Section 5.2 provides likelihood ratio test results supporting this claim.

3.3 Hypothesis testing

The PhyloDM model in (6) and (7) motivates a node-by-node testing strategy for cross-group comparison. To compare the proportion across G groups of observations, we carry out an MoM test using (3) individually for each A, i.e.

$$H_{0,A} : \boldsymbol{\pi}_A^{(1)} = \boldsymbol{\pi}_A^{(2)} = ... = \boldsymbol{\pi}_A^{(G)} \text{ vs } H_{a,A} : \text{otherwise}$$

Each of the MoM test statistic are calculated conditional on $\{N_i^{(g)}(A)|1 \le g \le G, 1 \le i \le n_g\}$, where $N_i^{(g)}(A)$ is the sum of OTU counts under A in the ith sample of gth group. The test statistic for $H_{0,A}$ only has degrees of freedom $(G-1)(k(A)-1)$, much smaller than the degrees of freedom for DM test as $(G-1)(K-1)$. The local DM test is therefore more powerful than the global DM test, provided that the extent of cross-group difference on the internal nodes is not diluted too much as we group multiple OTUs together. Obviously, the extent of dilution is largely determined by the tree structure. The ideal scenario is that OTUs placed under the same internal node A demonstrate increasing or decreasing abundance simultaneously for all samples in a certain group, so the test of $H_{0,R(A)}$ will be most effective. This also motivates using the phylogenetic tree to carry out the decomposition, as functionally similar OTUs tend to exhibit similar abundance changes within the same group.

The mean proportion of all OTUs across G groups are equal if and only if $H_{0,A}$ is true for all $A \in \mathcal{I}$. Therefore, we define the global null as $H_0 = \cap_{A\in\mathcal{I}} H_{0,A}$. Controlling the Type-I error on the global null is simply equivalent to controlling the family-wise error rate (FWER) across $H_{0,A}$'s.

The following theorem makes controlling FWER straightforward:

Theorem 1. *Let p_A be the MoM p-value for testing $H_{0,A}$. Under the global null $H_0 = \cap_{A\in\mathcal{I}} H_{0,A}$, p_A's are asymptotically mutually independent as the number of subjects in each group goes to infinity.*

A proof of this theorem is given in the appendix.

The independence of p-value under the null grants one of the following procedures to control the exact FWER at level α: (i) Sidak's procedure, in which one assigns equal Type I error $\alpha(A) = 1 - (1-\alpha)^{1/\mathcal{I}}$ to all A's (ii) allocate α_A according to the tree structure while constraining $1 - \prod_{A\in\mathcal{I}}(1-\alpha_A) = \alpha$. A reasonable choice is to put higher α_A on more ancient nodes where the occurrence of signals imply greater functional dissimilarities among groups. After choosing the individual Type I error thresholds, one can report the collection of nodes $\{A : p_A < \alpha_A, A \in \mathcal{I}\}$ as being significant.

4 Scan statistic over the tree tuples

Cross-group difference in distributions of taxa counts often occurs in clusters or chains on the phylogenetic tree. If one internal node exhibits significant difference in their relative proportion across several groups, then this is often associated with signals from at least one of its children or its parent. Figure 4 shows four examples of signal clusters on American Gut data using the top 100 OTUs with the highest counts. In each graph, subjects are divided into two groups according to different ingestion frequencies in one of the following diets: milk and cheese, seafood, sugary sweets and vegetable. (details in Section 5.1). The size of the circle on internal node A is proportional to -log(p_A) from the cross-group comparison (the circle colors are irrelevant here). It is apparent that large circles tend to form in chaining patterns, which prompts us to scan for signals in chains or clusters instead of on each node separately. Another motivation comes from the fact that the partitioning nature of the phylogenetic tree always leads to much smaller sample size on the bottom nodes (bottom means farthest from the root placed on top). Sharing information across nodes would alleviate the limitation to detect distributional differences on the bottom level.

Without prior knowledge of the length and shape of signal clusters, we only focus on triplets formulated by a certain internal node, its parent and one of its children. Each triplet has its own statistic defined as the sum of all the node statistics within, pooling signal strength from its members. The maximum of these statistics on all the triplets is then used to test the global null hypothesis. Our method belongs to the class of scan statistics [Glaz et al. (2001)], in which one searches for the signals over varying sizes of windows. In our case, each window denotes a particular branch of the phylogenetic tree. The shape of our designed triplet reflects our knowledge of correlated signals on the tree, while the size of the triplet achieves a compromise between signal pooling around neighboring nodes and the ability to detect alternatives in short chains. Since the exact distribution of the maximum statistic is unknown, we design a novel method to calculate

the upper and lower bound of its tail probability using low dimensional integrals that can be efficiently evaluated through standard numerical integration techniques. A decay rate of our approximation error with respect to the observed maximum statistic is also derived.

4.1 Overview

For each $A \in \mathcal{I}$ such that $R(A) \in \mathcal{I}$ and $\mathcal{C}(A) \cap \mathcal{I} \neq \emptyset$, we define a triplet to be the set of three consecutive internal nodes $\{A, R(A), \mathcal{C}(A)_i\}$ where $i \in \{1, 2, ..., k(A)\}$ satisfies $C(A)_i \in \mathcal{I}$. Let $\mathcal{B}$ be the set of all such triplets, and without loss of generality we write $\mathcal{B} = \{\mathcal{B}_1, \mathcal{B}_2, ..., \mathcal{B}_b\}$ where each $\mathcal{B}_i$ is a triplet and $b = |\mathcal{B}|$ depends on both K and the structure of the tree. We assume the ordering of elements in $\mathcal{B}$ obeys the following rule: $\{A, R(A), \mathcal{C}(A)_i\}$ always has a smaller index than (or appear in front of) $\{\tilde{A}, R(\tilde{A}), \mathcal{C}(\tilde{A})_j\}$ if $\tilde{A} \subset A$. Now we proceed to define the test statistic for $\mathcal{B}_i$ as follows. First, each of the p-values on the internal nodes can be inverted to a chi-square random variable with 1 degree of freedom, namely

$$Z_A = F_1^{-1}(p_A) \text{ for all } A \in \mathcal{I},$$

where F_j denotes the cumulative distribution function (CDF) of χ_j^2 distribution. Theorem 1 states that under the global null H_0, Z_A's are asymptotically mutually independent. In order to test the following hypothesis on each triplet

$$H_{0,\mathcal{B}_i} = \bigcap_{A \in \mathcal{B}_i} H_{0,A} \text{ vs } H_{a,\mathcal{B}_i} : \text{otherwise},$$

we define the statistic to be the sum of Z_A's within:

$$W_i = \sum_{A \in \mathcal{B}_i} Z_A \text{ for } i = 1, 2, ..., b. \tag{9}$$

It is apparent that each $W_i \sim \chi_3^2$ under $H_{0,\mathcal{B}_i}$. For the global null hypothesis $H_0 = \cap_{A \in \mathcal{I}} H_{0,A} = \cap_{i=1}^b H_{0,\mathcal{B}_i}$, we use the maximum of W_i's as the test statistic:

$$W = \max_{1 \leq i \leq b} W_i. \tag{10}$$

Since $\mathcal{B}_i$ overlaps with each other, W_i's are heavily correlated and the exact distribution of W is hard to derive. For hypothesis testing purposes, it suffices to only calculate the tail probability of W. Suppose our observed value of the maximum statistic is w, and let $B_i(w) = \{W_i > w\}$ be the event of ith triplet statistic exceeding w . Without incurring any confusion, we may drop w and simply write B_i. We are mainly interested in the global p-value $P(\bigcup_{i=1}^b B_i)$, which boils down to the problem of bounding the union probability.

The simplest upper bound of the union probability is the Bonferroni inequality:

$$P(\bigcup_{i=1}^{b} B_i) \leq \sum_{i=1}^{b} P(B_i)$$

There has been vast literature providing sharper bounds over the Bonferroni inequality in the past few decades. The results in Hunter (1976), Worsley (1982) and Efron (1997) suggest the following improvement:

$$\begin{aligned} P(\bigcup_{i=1}^{b} B_i) &= P(B_1) + P(B_2 \cap B_1^c) + P(B_3 \cap B_1^c \cap B_2^c) + ... \\ &\leq P(B_1) + \sum_{i=2}^{b} \min_{j<i} P(B_i \cap B_j^c) \end{aligned} \tag{11}$$

In particular, $\min_{j<i} P(B_i \cap B_j^c)$ is achieved at $j = i-1$ when the neighboring variables (W_{i-1}, W_i) have the highest pairwise correlation. Each of the term inside summation can be easily evaluated by numerical integration. It can be easily generalized to the union of more than two sets to improve approximation. More generally, the above inequality belongs to the class of approximations with the following representation:

$$P(\bigcup_{i=1}^{b} B_i) \leq \sum_{J \in S} (-1)^{|J|-1} f(J) P(\bigcap_{j \in J} B_j) \tag{12}$$

where $f(J)$ is some non-negative function on subset of $S = \{1, 2, ..., b\}$. Naiman and Wynn (1992) and Naiman and Wynn (1997) gave results regarding when (12) achieves equality. Following their work, Dohmen (2000) and Dohmen (2003) gave further improvement on the Bonferroni inequalities. There is also research on Bonferroni inequalities for particular applications, such as Dohmen and Tittmann (2004) on partition lattice and Taylor et al. (2007) on maxima over Gaussian random fields.

4.2 Bounding the union probability

Our upper bound of the union probability involves a decomposition of $\bigcup_{i=1}^{b} B_i$ into (i) a union of independent events and (ii) their complement in $\bigcup_{i=1}^{b} B_i$. The probability of the union of independent events can be exactly evaluated, while a similar strategy to (11) is applied to estimate its complement.

Specifically, let $\mathcal{M} = \{\mathcal{M}_1, \mathcal{M}_2, ..., \mathcal{M}_m\}$ be a class of disjoint nonempty subsets of $\mathcal{I}$ satisfying $\forall i \leq m, \exists j \leq b$ s.t $\mathcal{M}_i \subset \mathcal{B}_j$. For each i, define $M_i = \{\sum_{A \in \mathcal{M}_i} Z_A > w\}$ to be the exceeding event on $\mathcal{M}_i$. It follows that $\forall i \leq m, \exists j \leq b$ s.t $M_i \subset B_j$. Write $M = \bigcup_{i=1}^{m} M_i$. This leads to

$$P(\bigcup_{i=1}^{b} B_i) = P(M) + P(M^c) \cdot P(\bigcup_{i=1}^{b} B_i | M^c) \tag{13}$$

since $M \subset \bigcup_{i=1}^{b} B_i$. The independence of M_i's leads to a straightforward calculation of $P(M)$ as $P(M) = 1 - F_1(w)^{t_1} F_2(w)^{t_2} F_3(w)^{t_3}$ where $t_l = |\{i \leq m : |\mathcal{M}_i| = l\}|$ and $F_i(\cdot)$ is the CDF of χ^2_i distribution. Next, we approximate $P(\bigcup_{i=1}^{b} B_i | M^c)$ using a similar strategy to (11). It is apparent that enlarging M will always decrease $P(\bigcup_{i=1}^{b} B_i \cap M^c)$ and most likely the error of its upper bound, which makes our strategy superior to directly applying (11) to the B_i's.

The next question is how to choose an M as large as possible. An obvious optimality condition is that $\bigcup_{i=1}^{m} \mathcal{M}_i = \mathcal{I}$, because otherwise we can always enlarge M by $\{Z_A > w\}$ for a certain $A \in \mathcal{I} \setminus \bigcup_{i=1}^{m} \mathcal{M}_i$. Moreover, the elements in $\mathcal{M}$ should not be able to combine together and still belong to a certain element in $\mathcal{B}$, i.e. $\forall i_1, i_2 \leq m, \nexists j \leq b$ s.t. $\mathcal{M}_{i_1} \cup \mathcal{M}_{i_2} \subset \mathcal{B}_j$. This is because merging $\mathcal{M}_{i_1}$ and $\mathcal{M}_{i_2}$ enlarges M, i.e. $M_{i_1} \cup M_{i_2} \subset \{\sum_{A \in \mathcal{M}_{i_1} \cup \mathcal{M}_{i_2}} Z_A > w\}$. Since an exhaustive search over all combinations is computationally infeasible for large trees, we propose the following greedy algorithm that satisfies these optimality conditions:

(a) Order the elements in $\mathcal{I}$ as $A_1, A_2, ..., A_{|\mathcal{I}|}$ such that each internal node always appears in front of its children.

(b) Set $\mathcal{M} = \emptyset$ and $i = 1$.

(c) For each $i = 1, 2, ..., |\mathcal{I}|$, sequentially go through the following steps:

  (i) If $\exists j \leq m$ s.t. $A_i \in \mathcal{M}_j$, set $i \leftarrow i + 1$ and go back to (c).

  (ii) If $\exists j_1, j_2$ s.t. $\mathcal{C}(A_i)_{j_1} \in \mathcal{I}$ and $\mathcal{C}\big(\mathcal{C}(A_i)_{j_1}\big)_{j_2} \in \mathcal{I}$, set $\mathcal{M} \leftarrow \mathcal{M} \cup \{A_i, \mathcal{C}(A_i)_{j_1}, \mathcal{C}\big(\mathcal{C}(A_i)_{j_1}\big)_{j_2}\}$ and $i \leftarrow i + 1$. Go back to (c).

  (iii) If $\exists j_1$ s.t. $\mathcal{C}(A_i)_{j_1} \in \mathcal{I}$, set $\mathcal{M} \leftarrow \mathcal{M} \cup \{A_i, \mathcal{C}(A_i)_{j_1}\}$ and $i \leftarrow i + 1$. Go back to (c).

  (iv) Otherwise, set $\mathcal{M} \leftarrow \mathcal{M} \cup \{A_i\}$ and $i \leftarrow i + 1$. Go back to (c).

The above greedy algorithm seeks to incorporate the longest chain (with maximum of 3 nodes) starting from A_i and use its descendants as subsequent nodes, if A_i has not been included in $\mathcal{M}$ so far. Since the parent node is always considered ahead of its children, the resulting $\mathcal{M}$ will always satisfy the two aforementioned optimality conditions. As the algorithm prioritizes longer chains at each step, it effectively produces a large M that yields relatively accurate estimates of the union probability for our applications (numerical results to be shown later).

Figure 2 shows an example of $\mathcal{M}$ on a simple phylogenetic tree with $K = 13$ OTUs. Each internal node in $\mathcal{M}_i$ is assigned the same number i for $i = 1, 2, ..., 7$.

The remaining task is to put an upper bound on $P(\bigcup_{i=1}^{b} B_i | M^c)$. For each $i \leq b$, let $\mathcal{N}_i = \{j : |\mathcal{B}_j \cap \mathcal{B}_i| = 2 \text{ and } j < i\}$. Apparently $|\mathcal{N}_i| \leq 2$ for all i because of the ordering of $\mathcal{B}_i$'s. Write $B_{\mathcal{N}_i} = \bigcup_{j \in \mathcal{N}_i} B_j$ for short. Now we proceed as follows:

$$P(\bigcup_{i=1}^{b} B_i | M^c) \leq \sum_{i=1}^{b} P(B_i \cap B^c_{\mathcal{N}_i} | M^c) \tag{14}$$

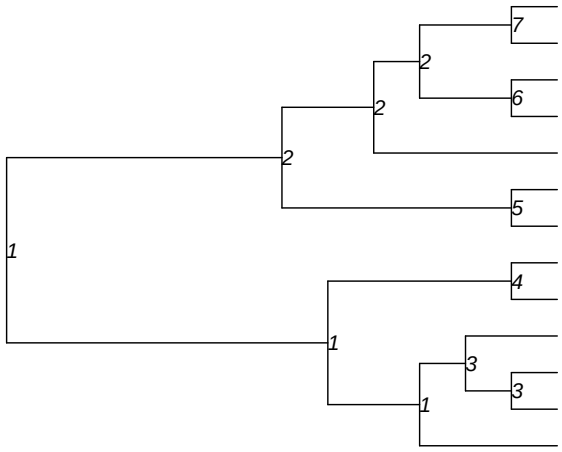

Figure 2: Example configuration of $\mathcal{M}$ using the greedy algorithm.

The equation in (14) is very similar to (11) in that for each B_i, it only includes the highest correlated events $B_{\mathcal{N}_i}$ in order to minimize the right side of the equation. It is worth noting that $P(B_i|M^c) = 0$ if $\mathcal{B}_i \in \mathcal{M}$, hence the strategy of prioritizing triplets while constructing $\mathcal{M}$. To efficiently evaluate each of the terms in the right side of (14), notice that the distributions of Z_A conditioned on M^c are the same as the product of a truncated chi-square and an independent Dirichlet random variable, so their density function can be expressed using chi-square CDFs. Let $f_i(\cdot)$ and $F_i(\cdot)$ denote the density and CDF of χ^2_i distribution respectively, then the marginal density of Z_A conditional on M^c becomes

$$f_A(z|M^c) = \begin{cases} \frac{f_1(z)}{F_1(w)}, \text{ if } |\mathcal{M}(A)| = 1 \wedge w \leq z \\ \frac{F_1(w-z)f_1(z)}{F_2(w)}, \text{ if } |\mathcal{M}(A)| = 2 \wedge w \leq z \\ \frac{F_2(w-z)f_1(z)}{F_3(w)}, \text{ if } |\mathcal{M}(A)| = 3 \wedge w \leq z \\ 0, \text{ otherwise} \end{cases} \tag{15}$$

where we define $\mathcal{M}(A) = \mathcal{M}_i$ if $A \in \mathcal{M}_i$. The existence and uniqueness of $\mathcal{M}(A)$ is guaranteed by the fact that $\{\mathcal{M}_1, \mathcal{M}_2, ..., \mathcal{M}_m\}$ are disjoint with $\bigcup_{i=1}^m \mathcal{M}_i = \mathcal{I}$.

The joint density of Z_{A_1} and Z_{A_2} for $\forall A_1, A_2 \in \mathcal{I}$ is

$$f_{A_1,A_2}(z_1, z_2|M^c) = \begin{cases} f_{A_1}(z_1|M^c) f_{A_2}(z_2|M^c), \text{ if } \mathcal{M}(A_1) \cap \mathcal{M}(A_2) = \emptyset \\ \frac{\prod_{i=1}^2 f_1(z_i)}{F_2(w)}, \text{ if } \mathcal{M}(A_1) = \mathcal{M}(A_2) \wedge |\mathcal{M}(A_1)| = 2 \wedge \sum_{i=1}^2 z_i \leq w \\ \frac{F_1(w - \sum_{i=1}^2 z_i) \prod_{i=1}^2 f_1(z_i)}{F_3(w)}, \text{ if } \mathcal{M}(A_1) = \mathcal{M}(A_2) \wedge |\mathcal{M}(A_1)| = 3 \wedge \sum_{i=1}^2 z_i \leq w \\ 0, \text{ otherwise} \end{cases} \tag{16}$$

Given w, we pre-calculate the density functions in (15) and (16), as well as the CDFs of Z_A using (15) and of $Z_{A_1} + Z_{A_2}$ using (16) up to a certain precision and store them into the memory. This turns each term in the right side of (14) into at most four-dimensional integrals. We further evaluate these integrals using the functions `cuhre` and `suave` in R package `R2Cuba` [Hahn (2005)].

Plugging (14) into (13) gives

$$P_0 = P(\bigcup_{i=1}^{b} B_i) \leq P_U = P(M) + P(M^c) \cdot \sum_{i=1}^{b} P(B_i \cap B^c_{\mathcal{N}_i} | M^c), \tag{17}$$

where P_0 is the actual p-value and P_U is its upper bound. Let $\epsilon_U = P_U - P_0$ be the error of our approximation. Using a similar strategy to Theorem A1 in Taylor et al. (2007), it follows that

$$\begin{aligned}
\epsilon_U &= P(M^c) \sum_{i=1}^{b} \big(P(B_i \cap B^c_{\mathcal{N}_i} | M^c) - P(B_i \cap B^c_{i-1} \cap B^c_{i-2} \cap ... B^c_1 | M^c) \big) \\
&= P(M^c) \sum_{i=1}^{b} P(B_i \cap B^c_{\mathcal{N}_i} \cap \big(\bigcup_{j<i, j \notin \mathcal{N}_i} B_j \big) | M^c) \\
&\leq P(M^c) \sum_{i=1}^{b} P(\bigcup_{j<i, j \notin \mathcal{N}_i} (B_i \cap B_j) | M^c) \\
&\leq P(M^c) \sum_{i=1}^{b} \sum_{j<i, j \notin \mathcal{N}_i} P(B_i \cap B_j | M^c)
\end{aligned} \tag{18}$$

Each term in (18) can be evaluated by at most three dimensional numerical integral using the pre-calculated densities and CDFs. This also establishes

$$P_0 \in \big(P_U - P(M^c) \sum_{i=1}^{b} \sum_{j<i, j \notin \mathcal{N}_i} P(B_i \cap B_j | M^c), P_U \big) \tag{19}$$

In the next section we give the numerical results of P_U and upper bound of ϵ_U using the phylogenetic tree from the American Gut dataset. In addition, we have the following theorem on the convergence rate of the relative error with regards to the observed statistic w.

Theorem 2. *Given the set of all triplets $\mathcal{B}$ and the partition $\mathcal{M}$ on the internal nodes $\mathcal{I}$, define the following quantities:*

- $\xi_1 = |\{(i,j) : \mathcal{B}_i \cap \mathcal{B}_j = \emptyset, \mathcal{B}_i \notin \mathcal{M}, \mathcal{B}_j \notin \mathcal{M}$ *and* $1 \leq j < i \leq b\}|$
- $\xi_2 = |\{(i,j) : |\mathcal{B}_i \cap \mathcal{B}_j| = 1, \mathcal{B}_i \notin \mathcal{M}, \mathcal{B}_j \notin \mathcal{M}$ *and* $1 \leq j < i \leq b\}|$
- $\xi_3 = |\{i : |\mathcal{M}_i| = 3$ *and* $1 \leq i \leq m\}|$

Then under the condition that $(\xi_3 - 1)\big(1 - F_3(w_T)\big) < 0.1$ *and* $w_T \geq 12$*, we have*

$$\begin{aligned}\frac{\epsilon_U}{P_U} &< \frac{\xi_1}{0.95\xi_3 + \xi_T}\big(1 - F_3(w)\big) + \frac{0.9\xi_2}{0.95\xi_3 + \xi_T}\sqrt{\frac{\pi}{2}} \cdot \frac{1}{w} \\ &= \mathcal{O}(e^{-\frac{w}{6}}) + \mathcal{O}(w^{-1})\end{aligned}$$

for all $w \geq w_T$*, where*

$$\xi_T = \frac{\sum_{i=1}^{b} P(B_i \cap B_{\mathcal{N}_i}^c \cap M^c)}{1 - F_3(w)} \text{ is evaluated at } w = w_T$$

The proof is presented in the appendix.

4.3 Comparison with Monte-Carlo simulation

In this section we compare the lower and upper bound in (19) with Monte-Carlo simulated p-values. Each round of simulation produces 5×10^4 simulated maximum triplet statistics, and we use their proportion of exceeding w as the estimated p-value. The maximum triplet statistic is simulated through generating the χ_1^2 distributed Z_A's for all $A \in \mathcal{I}$ and then applying (9) and (10). We draw the comparison for a variety of scenarios with different numbers of OTU $K \in \{50, 75, 100\}$ and different thresholds $w \in \{15, 20, 25\}$. Given K, the tree structure is obtained from keeping the top K OTUs with the highest count in all feces samples from the American Gut dataset (to be introduced later). In each scenario, we provide the histogram of simulated p-values over 500 rounds.

Figure 3 shows that our bounds consistently contain the center of the simulated p-values. Since the Monte Carlo p-values are merely binomial proportions, the ratio of their spread (measured by standard deviation) to P_0 goes to infinity as $P_0 \to 0$. In contrast, our method gives a ratio that tends to zero by Theorem 2. This makes our approach particularly useful for scenarios with a large number of tests where the p-value threshold is very small after multiple testing correction. In order to achieve a constant relative error, the computation time of Monte Carlo method has to scale up as the desired p-value goes down (assuming the tree size is fixed), whereas our method has much slower increment of time cost due to convergence of numerical integration.

5 Application to American Gut Project

American Gut Project [McDonald et al. (2015a)] is an open-access and crowd-sourced initiative that involves the public into the research of human microbiome and aims at providing a much more comprehensive reference set than the previous Human Microbiome Project [Human Microbiome Project Consortium (2012)]. After contributing to the project fund, participants complete a questionnaire and ship their microbiome sample to the sequencing lab currently located at University of California, San Diego. The questionnaire covers a wide range of topics regarding demographic information, diet, lifestyle,

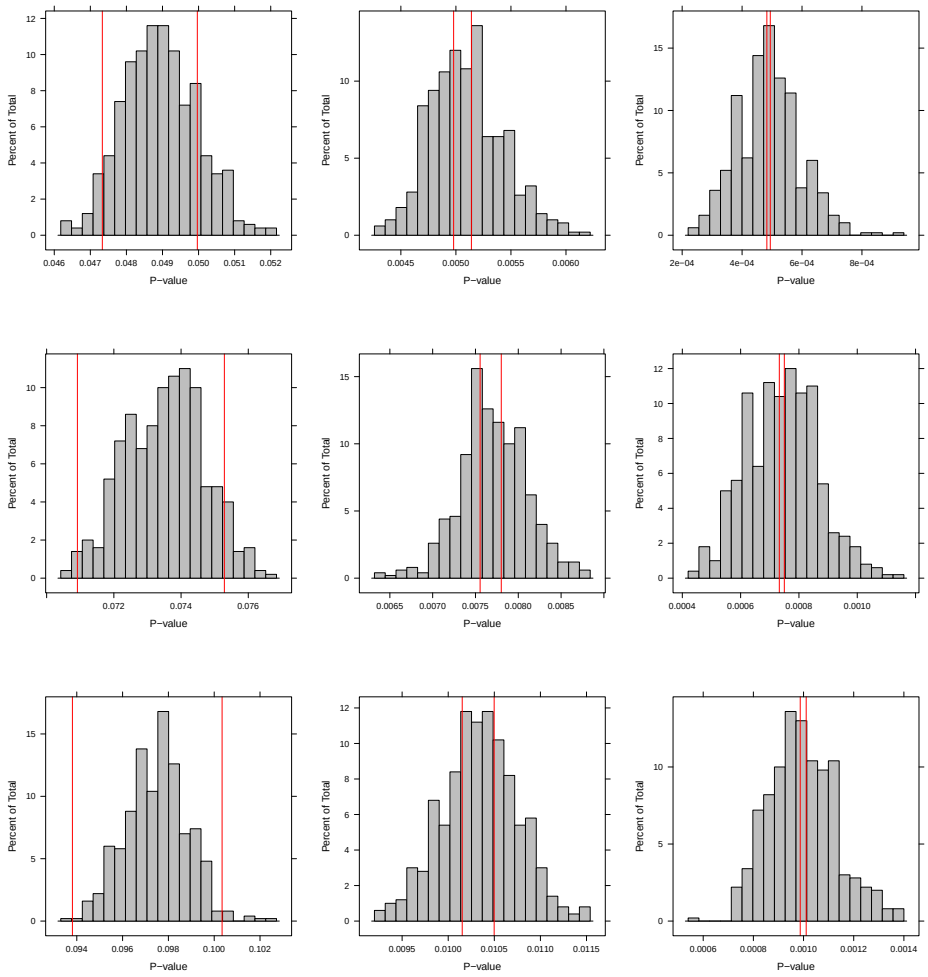

Figure 3: Comparison between the interval bound and simulated p-values. Each simulated p-value is the proportion exceeding w over 5×10^4 runs. Red lines indicate the upper and lower bound as in (19). Top row: $K = 50$, middle row: $K = 75$, and bottom row: $K = 100$. Left column: $w = 15$, middle column: $w = 20$, and right column $w = 25$.

etc. Sampling sites include skin, tongue and feces, although the vast majority of participants provided the feces sample. The samples are sequenced on 16s rRNA and further processed by QIIME [Caporaso et al. (2010)] pipeline to produce the OTUs and the phylogenetic tree. The 2016 May 16 cohort of public dataset includes more than eight thousands of subjects, with median of sequences per individual as 14680 and standard deviation as 32455.

5.1 Cross-group comparison

Our focus is comparison of the feces microbiome across different diet habits. We pick the top 100 OTUs with the highest count summing over all feces samples. The phylogenetic

tree on these OTUs is fully binary. We also select a total of seven categories of diet from the questionnaire. Each diet divides the samples into two groups; group 1 consists of individuals with ingestion rate less than three times per week, and group 2 corresponds to more than or equal to three times per week. Since the questions are not compulsory, a large number of subjects do not leave any response. The diet names and their sample sizes in both groups are as follows: fermented plant (902 vs 3081), fruit (2393 vs 1685), milk and cheese (1782 vs 2304), poultry (1446 vs 2668), seafood (568 vs 3521), sugary sweet (1574 vs 2543) and vegetable (3496 vs 584).

For each diet type, we fit the PhyloDM model and use the technique in Section 3.3 to test equality of proportions between two groups. Table 1 presents the p-values using PhyloDM and DM on the 100 OTUs. The PhyloDM column contains P_U derived in Section 4.2 and the upper bound of its error ϵ_U in the parenthesis. We also provide DM p-values after grouping the 100 OTUs into family, order and class levels, respectively. The grouping operation based on taxonomy is a common practice in recent papers including La Rosa et al. (2012) and Chen and Li (2013). At each taxonomic level, all OTUs with missing taxa information are placed into the same group. This leads to a total of 22 categories on family level, 11 categories on order level and 9 categories on class level. The DM p-values are calculated using the R package `HMP`.

All diet habit comparisons exhibit significant PhyloDM p-values at 0.05 level except fermented plant. This is consistent with the findings in Turnbaugh et al. (2009) and David et al. (2014), both of which established that the human gut microbiome is highly sensitive to the dietary nutrient composition. The p-values of DM on OTUs, on the other hand, fail to reach significance for fermented plant, poultry and sugary sweet. DM p-values on family, order or class levels fail to reach significance in even more comparisons. Notice that here we only provide the raw DM p-values. If we correct for multiple testing on these 4 different levels, DM significance will further diminish. Combining OTUs into family, order and class level also give contradictory results in certain cases. For example, the fermented plant comparison has decreasing DM p-values as the taxonomic rank goes up, but the pattern entirely reverses in fruit comparison.

We further visualize the significant internal nodes in Figure 4 for four of the diet comparisons: milk and cheese, seafood, sugary sweets and vegetable. Using a simple binary search, we find that $w = 16.579$ yields $P_U = 0.05$ with $\epsilon_U \leq 2.48 \times 10^{-3}$. All triplets with test statistics greater than the above threshold, i.e. $W_i > 16.579$, are plotted in red. In some cases the triplets overlap with each other, leading to a much longer chain than the original setup. We also provide the taxonomy for all internal nodes that belong to a certain significant triplet in Table 2. The internal node taxon is defined according to the following algorithm: starting from kingdom, repeatedly decrease the rank by one level until the descendant OTUs of that particular internal node no longer share the same taxa on the next lower rank (missing taxa on OTUs are excluded). The algorithm then picks the common taxon of the descendant OTUs on the rank at which the algorithm stops. In other words, the internal node taxon reflects the finest classification upon which all of its descendant OTUs agree. This illustrates another advantage of PhyloDM, as it is capable of reporting significant taxa along the entire spectrum of taxonomic ranks.

		DM p-values			
Diet	**PhyloDM**	**OTU**	**Family**	**Order**	**Class**
Fermented plant	0.308 (0.036)	0.377	0.147	0.073	0.038
Fruit	8.67×10^{-5} (1.78×10^{-6})	2.81×10^{-3}	0.012	0.274	0.218
Milk and cheese	1.06×10^{-4} (2.21×10^{-6})	0.029	0.262	0.162	0.285
Poultry	0.023 (9.00×10^{-4})	0.158	0.287	0.744	0.691
Seafood	6.77×10^{-5} (1.40×10^{-6})	1.75×10^{-4}	0.194	0.344	0.772
Sugary sweet	5.12×10^{-3} (1.49×10^{-4})	0.719	0.558	0.478	0.815
Vegetable	7.31×10^{-5} (1.51×10^{-6})	3.77×10^{-3}	1.88×10^{-3}	0.014	0.014

Table 1: DM and PhyloDM p-values for testing microbiome compositions across different diet habits. For PhyloDM, we provide P_U and the upper bound on ϵ_U shown in parenthesis. For DM, we provide p-values directly on the 100 OTUs as well as after grouping the 100 OTUs into family, order and class levels, respectively.

5.2 DM vs PhyloDM test

We can also test the model fit of PhyloDM against DM directly on the OTUs. Since Section 3.2 shows that DM is nested in the PhyloDM family, we can use the likelihood ratio test (LRT) for

$$H_0 : \exists \nu > 0 \text{ s.t. } \forall A \in \mathcal{I}, \nu_A = \nu \sum_{\omega \in A} \pi_\omega \text{ (DM) vs } H_a : \text{otherwise (PhyloDM)}$$

with the test statistic defined as

$$\Lambda(\boldsymbol{x}) = -2 \log \frac{\mathcal{L}(\hat{\nu}, \hat{\boldsymbol{\pi}})}{\mathcal{L}_T(\{(\hat{\nu}_A, \hat{\boldsymbol{\pi}}_A) : A \in \mathcal{I}\})} \sim \chi^2_{|\mathcal{I}|-1} \text{ under } H_0, \tag{20}$$

where $(\hat{\nu}, \hat{\boldsymbol{\pi}})$ in the numerator of (20) are MLEs of the DM model, and each $(\hat{\nu}_A, \hat{\boldsymbol{\pi}}_A)$ in the denominator are obtained through maximizing the PhyloDM conditional likelihood (5). We use the `dirmult` package in R to calculate the MLE estimates by Newton-Raphson method. When the algorithm fails to converge on certain internal nodes, we switch to the low-storage BFGS optimization implemented in package `nloptr`. The degrees of freedom in (20) is $|\mathcal{I}| - 1$ for a binary phylogenetic tree since (i) $\dim(\boldsymbol{\pi}) = \dim(\{\boldsymbol{\pi_A} : A \in \mathcal{I}\}) = K - 1$, and (ii) $\dim(\{\nu_A : A \in \mathcal{I}\}) = |\mathcal{I}|$.

Table 3 shows the LRT result. The test is separately applied to male and female Caucasians living in a variety of geographic regions. Each region is consisted of certain

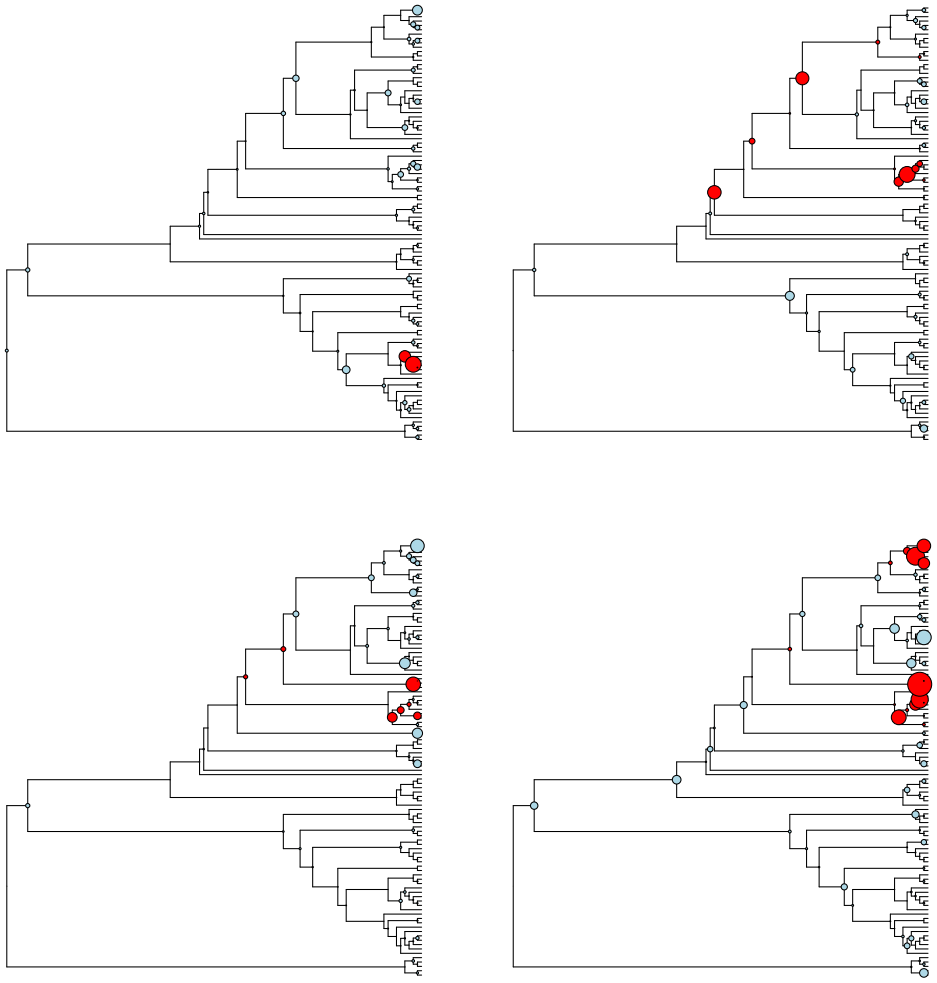

Figure 4: Significant triplets from PhyloDM testing. Top left: Milk and cheese, top right: seafood, bottom left: sugary sweets and bottom right: vegetable. The size of the circle on internal node A is proportional to $-\log(p_A)$. Triplets with their sum statistic $W_i > 16.579$ are plotted in red, indicating their significance.

states in the U.S. defined according to Bureau of Economic Analysis. The degree of freedom for all tests is 98 since our phylogenetic tree is binary and $|\mathcal{I}| = |K| - 1 = 99$. All scenarios yield LRT p-values less than 10^{-10}, which indicates significantly improved fit on the data using PhyloDM. We also note that $\Lambda(x)$ in general increases with the sample size, as evidence towards heterogeneity in OTU dispersion strengthens with more available data.

Diet	Phylum	Order	Family	Genus
Fermented plant	—	—	—	—
Fruit	Firmicutes	Clostridiales	Clostridiaceae	Clostridium
	—	—	Ruminococcaceae	Faecalibacterium
Milk and cheese	—	—	—	Bacteroides
Poultry	—	Clostridiales	—	Coprococcus
Seafood	Firmicutes	Clostridiales	Lachnospiraceae	Coprococcus
	—	—	Ruminococcaceae	Ruminococcus
Sugary sweet	Firmicutes	Clostridiales	Lachnospiraceae	Coprococcus
	—	—	Ruminococcaceae	—
Vegetable	Firmicutes	Clostridiales	Lachnospiraceae	Blautia
	—	—	Ruminococcaceae	Coprococcus
	—	—	—	Lachnospira

Table 2: Taxa on significant triplets from PhyloDM hypothesis testing for each diet comparison. Each internal node that belongs to a certain significant triplet is assigned a taxon based on its descendant OTUs (details described in section 5.1). Only the lowest level taxon is reported for each internal node. We omit the class rank since there are no significant internal nodes on such level in any cross-group comparisons.

5.3 Simulation

We use two simulation strategies to evaluate the power of PhyloDM vs DM under various conditions. From all subjects with feces samples in American Gut dataset, we extract a total of 670 individuals who identified themselves as male Caucasian living in the far west (including the states of Alaska, California, Hawaii, Nevada, Oregon and Washington). In each round of simulation, these selected samples are randomly divided into two equal-sized groups to generate data under the global null. For data under the alternative, the first simulation strategy randomly selects an OTU and increases its count by a fixed percentage for all samples in the second group, whereas the second simulation strategy random selects an internal node and increases the count all of its descendant OTUs equally by a fixed percentage for all samples in the second group. We use the same 100 OTUs as before and produce 5000 rounds of simulation.

Figure 5 demonstrates the distribution of DM and PhyloDM p-values under the global null. We fit three separate DM on 100 OTUs, family level and class level. The histogram of PhyloDM p-values is produced by using all the p-values on the internal nodes. Surprisingly, the distribution of p-values for DM on the OTUs is far from being uniform on (0,1), which leads to overly conservative estimates and much less power. The discrepancy alleviates as we group more OTUs into family or class level, although its empirical distribution is still noticeably skewed. This phenomenon reflects the fact that DM is severely under-parametrized for microbiome data even in low dimensions, as it fits a single disper-

	Male		Female	
Region	$\boldsymbol{\Lambda(x)}$	**Sample size**	$\boldsymbol{\Lambda(x)}$	**Sample size**
Far West	14179.76	671	15768.10	786
Great Lakes	4025.36	183	5541.45	280
Mideast	7497.80	331	8112.80	398
New England	5630.21	244	5793.38	274
Rock Mountain	6084.90	245	6676.25	301
Southeast	7030.43	327	7383.72	371
Southwest	3442.69	155	3675.52	190

Table 3: Likelihood ratio test for DM vs PhyloDM. The test is separately applied to male and female Caucasians in a variety of geographic regions. Each test is accompanied by the LRT statistic $\Lambda(x)$ and the sample size.

sion parameter that simultaneously controls all categories. In contrast, PhyloDM solves this issue through fitting a family of dispersion parameters $\{\nu_A : A \in \mathcal{I}\}$ that leads to better calibrated p-values.

Figure 6 and 7 show the ROC and power curves when we use the first simulation strategy (increase the OTU count) to increase the count of a random OTU. We provide the result for (i) DM on the OTUs (ii) PhyloDM using the maximum of the single node statistic, or $\max_{A\in\mathcal{I}} Z_A$ and (iii) PhyloDM using the maximum of the triplet statistic, or $\max_{i\leq b} W_i$. Both PhyloDM methods give improved performance compared to DM due to highly localized signal.

Figure 8 and 9 show the ROC and power curves when we use the second simulation strategy to increase the count of all OTUs under a random internal node. The minimum number of OTUs under the randomly selected internal node controls degree of localization in the signal. This simulation setup reflects the more biologically meaningful scenario in which a number of taxa exhibit differences in the between-group comparison. In all cases, PhyloDM consistently provides higher power than DM. The PhyloDM 3-node method also provides higher power than PhyloDM 1-node at moderate increment levels. When the increment level is high, there will be a certain Z_A whose value dominates all other node statistics, so the extra gain from pooling signal strength within triplets diminishes.

6 Discussion

Traditional DM models are less than desirable for microbiome data as they ignore species relatedness, have limited power to detect realistic signals, and yield results that are not easy to interpret. Our proposed PhyloDM handles the above issues through a cascade of local DMs with varying degrees of resolutions on the phylogenetic tree. We further take advantage of the correlated signals on the tree through a scan statistic approach and provide upper and lower bound on its tail probability for testing cross-group differences. Both empirical results on American Gut data and simulation study demonstrated that

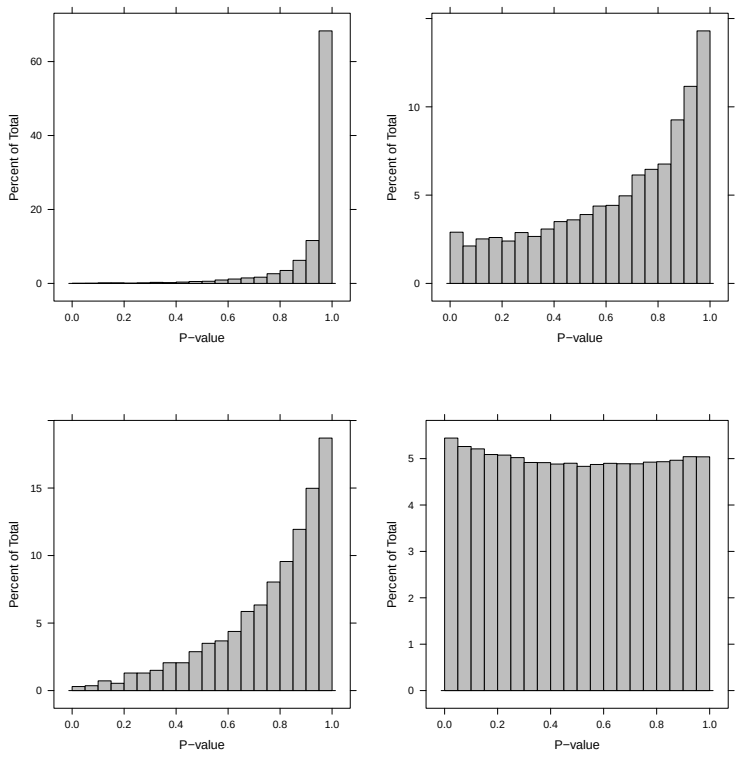

Figure 5: P-value histograms under the global null. We randomly place the 670 samples from Caucasian male living in far west into two equal-sized groups and produce their p-values for 5000 rounds. Top left is DM on the OTUs, top row right is DM on family level, bottom left is DM on class level, and bottom right is PhyloDM.

our methods give promising results.

The PhyloDM model we developed is a generalization of DM with $|\mathcal{I}| - 1$ more dispersion parameters. It is an interesting question whether one could stepwise tune the model (hence the number of parameters) from DM to PhyloDM. To start, the PhyloDM representation in (4) shrinks to the degenerate DM if $\exists \nu > 0$ s.t. $\nu_A = \nu \sum_{w \in A} \pi_w$ for all $A \in \mathcal{I}$. This condition is equivalent to $\nu_A = \nu_{R(A)} \pi_{R(A),i}$ for $\forall A \in \mathcal{I} \backslash \{\Omega\}$ with $A = \mathcal{C}(R(A))_i$. Stepwise tuning can be achieved through only requiring $\nu_A = \nu_{R(A)} \pi_{R(A),i}$ to hold over $A \in \tilde{\mathcal{I}}$ where $\tilde{\mathcal{I}} \subset \mathcal{I} \backslash \{\Omega\}$ controls the effective degrees of freedom. Apparently $\tilde{\mathcal{I}} = \emptyset$ leads to PhyloDM and $\tilde{\mathcal{I}} = \mathcal{I} \backslash \{\Omega\}$ leads to DM, so any choice of $\tilde{\mathcal{I}}$ in the middle yields a model between the two extremes. Standard model selection techniques such as information criterion or cross validation can then be applied to identify the best choice. Although the existence of such spectrum grants substantial flexibility, we note that it can be computationally infeasible to examine the model fit of all $2^{|\mathcal{I}|-1}$ possible configurations. A potential workaround is to enlarge $\mathcal{T}$ stepwise by a greedy algorithm or use dynamic programming, but it is not clear under which conditions we are guaranteed to recover the global optimum.

The PhyloDM framework can also be adopted to incorporate continuous variable of

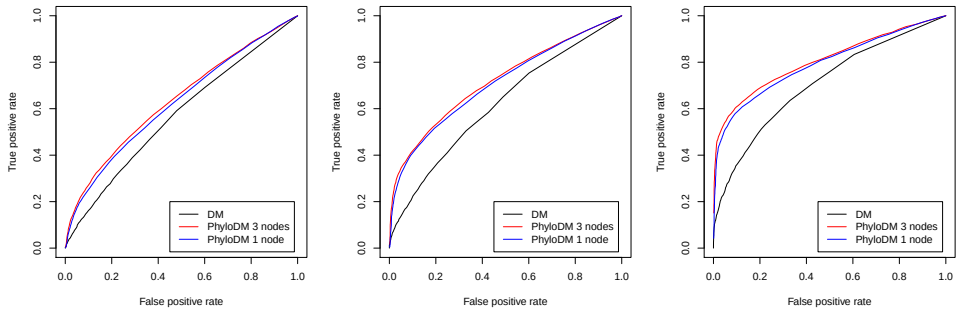

Figure 6: ROC curves from increasing the count of a random OTU. For left to right, the percentage increment is set as 100%, 150% and 250%

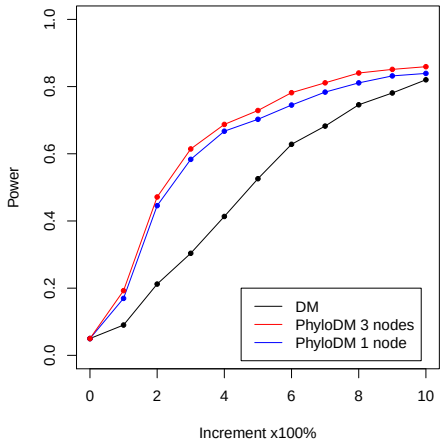

Figure 7: Power of DM and PhyloDM with regard to different increment in a random OTU at false positive rate = 0.05.

interest and to adjust for the effects of confounders. When the tree is fully binary, we let $\lambda_A = \pi_{A,1}$ to fully represent $\boldsymbol{\pi}_A = (\pi_{A,1}, 1 - \pi_{A,1})$ in (4). Then we can build separate logistic regression models for each A:

$$\log \frac{\lambda_A}{1 - \lambda_A} = \beta_{A,0} + \beta_{A,1} u + \sum_{i=1}^{s} \beta_{A,s+1} c_s \tag{21}$$

where $\beta_{A,i}$ is the ith regression coefficient, u denotes the continuous variable of interest and c_1, c_2, ..., c_s are the confounders. After obtaining maximum likelihood estimates of the coefficients as well as ν_A, we test the significance of u's coefficient to produce a p-value and use the same scan statistics approach in Section 4 to borrow strength from neighboring nodes. One could also apply DM penalized regression [Chen and Li (2013)] individually for each node, although the lack of p-value prevents the usage of scan statistic. Further studies are needed to design penalized regression strategies while jointly considering all nodes.

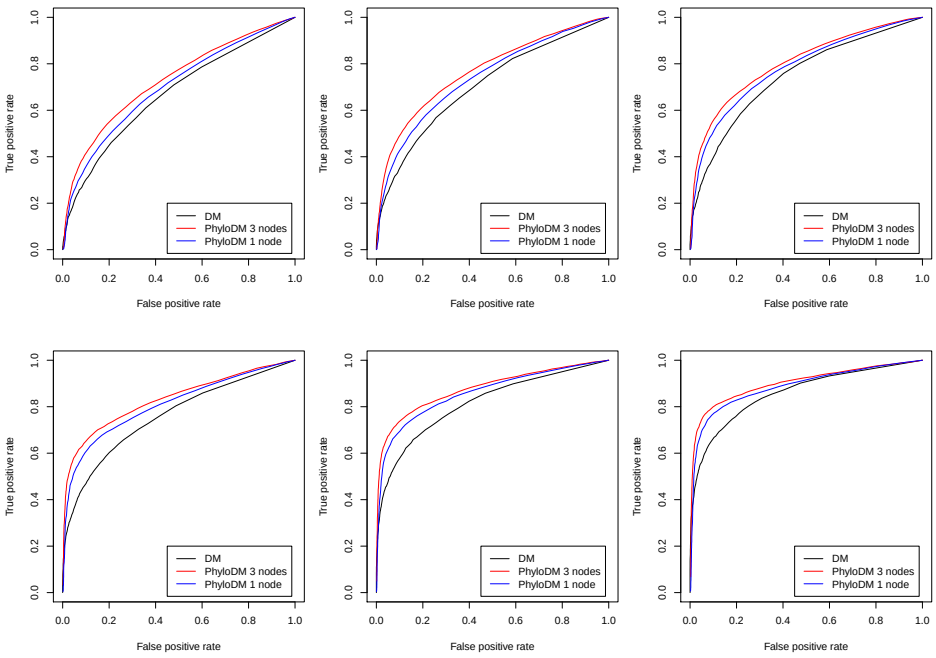

Figure 8: ROC curves from increasing the count of all OTUs under a random internal node. The top row and the bottom row have the percentage increment set as 50% and 75% respectively. From left to right column, the minimum number of OTUs under the chosen internal node is 2, 3 and 5.

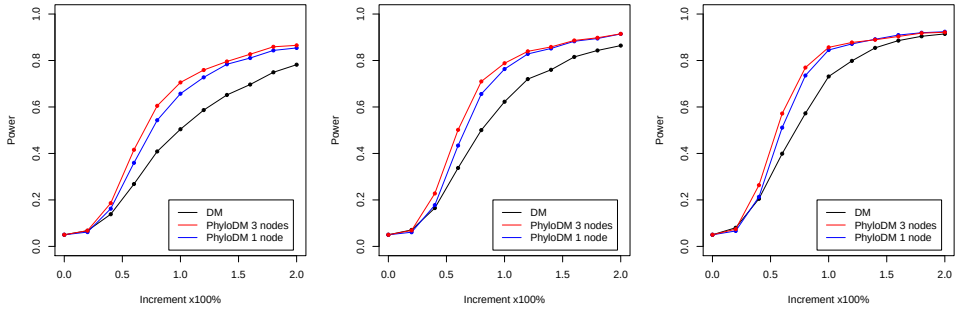

Figure 9: Power of DM and PhyloDM with regard to different increment in all OTUs under a random internal node at false positive rate = 0.05. From left to right, the minimum number of OTUs under the chosen internal node is 2, 3 and 5.

Acknowledgements

L. Ma is partly supported by NSF grant DMS-1612889 and a Google Faculty Research Award. D. L. Nicolae is supported in part by NSF grants R01-MH101820 and R01-HL129735.

Appendix

A1. Proof of Theorem 1

The elements in $\mathcal{I}$ can be ordered as $A_1, A_2, ..., A_{|\mathcal{I}|}$ such that each parent node always appears in front of its children. Let p_{A_l} be the p-value for testing H_{0,A_l}. Without loss of generality, let us assume $A_{|\mathcal{I}|} = \{K-1, K\}$. For any subject with OTU counts $\boldsymbol{x}$, the probability function (6) can be written as

$$\begin{aligned} f_T(\boldsymbol{x}) &= \prod_{l=1}^{\mathcal{I}} f\big(\boldsymbol{x}(A_l)|N(A_l)\big) \\ &= \prod_{l=1}^{\mathcal{I}-1} f\big(\boldsymbol{x}(A_l)|N(A_l)\big) \cdot f\big(\boldsymbol{x}(A_{|\mathcal{I}|})|N(A_{|\mathcal{I}|})\big) \\ &= f_T\big(x_1, x_2, ..., x_{K-2}, N(A_{|\mathcal{I}|})\big) f\big(\boldsymbol{x}(A_{|\mathcal{I}|})|N(A_{|\mathcal{I}|})\big), \end{aligned}$$

which yields

$$f_T\big(\boldsymbol{x}|N(A_{|\mathcal{I}|})\big) = f_T\big(x_1, x_2, ..., x_{K-2}, N(A_{|\mathcal{I}|})|N(A_{|\mathcal{I}|})\big) f\big(\boldsymbol{x}(A_{|\mathcal{I}|})|N(A_{|\mathcal{I}|})\big).$$

Since the above conditional independence relationship holds for all subjects, it follows that $p_{A_{|\mathcal{I}|}}$ is independent of all other p_{A_l}'s conditional on $N(A_{|\mathcal{I}|})$. Therefore under H_0,

$$\begin{aligned} P(\bigcup_{l=1}^{|\mathcal{I}|}\{p_{A_l} \le \alpha_l\}) &= E\Big(P(\bigcup_{l=1}^{|\mathcal{I}|}\{p_{A_l} \le \alpha_l\}|N(A_{|\mathcal{I}|}))\Big) \\ &= E\Big(P(\bigcup_{l=1}^{|\mathcal{I}|-1}\{p_{A_l} \le \alpha_l\}|N(A_{|\mathcal{I}|}))P\big(p_{A_{|\mathcal{I}|}} \le \alpha_{|\mathcal{I}|}|N(A_{|\mathcal{I}|})\big)\Big) \\ &= \alpha_{|\mathcal{I}|} E\Big(P(\bigcup_{l=1}^{|\mathcal{I}|-1}\{p_{A_l} \le \alpha_l\}|N(A_{|\mathcal{I}|}))\Big) \\ &= \alpha_{|\mathcal{I}|} P(\bigcup_{l=1}^{|\mathcal{I}|-1}\{p_{A_l} \le \alpha_l\}). \end{aligned}$$

where the second last equation requires the asymptotic distribution of (3) holds so that $P\big(p_{A_{|\mathcal{I}|}} \le \alpha_{|\mathcal{I}|}|N(A_{|\mathcal{I}|})\big) = \alpha_{|\mathcal{I}|}$

Repeating the above procedures iteratively for $A_{|\mathcal{I}|-1}$, $A_{|\mathcal{I}|-2}$, ... A_1 gives

$$P(\bigcup_{l=1}^{|\mathcal{I}|}\{p_{A_l} \le \alpha_l\}) = \prod_{l=1}^{|\mathcal{I}|}\alpha_l.$$

A2. Proof of Theorem 2

By (18),

$$\begin{aligned}\epsilon_U &\le \sum_{i=1}^{b} \sum_{j<i, j\notin\mathcal{N}_i} P(B_i \cap B_j \cap M^c) \\ &= \sum_{i\le b, \mathcal{B}_i\notin\mathcal{M}} \sum_{j<i, j\notin\mathcal{N}_i, \mathcal{B}_j\notin\mathcal{M}} P(B_i \cap B_j \cap M^c).\end{aligned}$$

The elements inside the summation sign above fall in one of the following two categories

(i) $|\mathcal{B}_i \cap \mathcal{B}_j| = 0$ which means $P(B_i \cap B_j \cap M^c) < P(B_i \cap B_j) = \big(1 - F_3(w)\big)^2$

(ii) $|\mathcal{B}_i \cap \mathcal{B}_j| = 1$ so that $P(B_i \cap B_j \cap M^c) \le P(\sum_{i=1}^{3} Y_i > w, \sum_{i=3}^{5} Y_i > w, Y_i < w$ for all $i)$, where Y_i's are i.i.d. chi-square distributed with 1 degree of freedom.

Let $f_1(y) = \frac{y^{-\frac{1}{2}} e^{-\frac{y}{2}}}{\sqrt{2\pi}}$ be the density function of χ_1^2. Conditioning on Y_3 gives the following upper bound on category (ii):

$$\begin{aligned}P(\sum_{i=1}^{3} Y_i > w, \sum_{i=3}^{5} Y_i > w, Y_i < w) &< \int_0^w f_1(y) P(\sum_{i=1}^{3} Y_i > w | Y_3 = y) \cdot \\ &\quad P(\sum_{i=3}^{5} Y_i > w | Y_3 = y) dy \\ &= \int_0^w \frac{y^{-\frac{1}{2}} e^{-\frac{y}{2}}}{\sqrt{2\pi}} \big(1 - F_2(w-y)\big)^2 dy \\ &= \frac{1}{\sqrt{2\pi}} \int_0^w y^{-\frac{1}{2}} e^{-\frac{y}{2}} e^{-(w-y)} dy \\ &= \frac{e^{-w}}{\sqrt{2\pi}} \int_0^w y^{-\frac{1}{2}} e^{\frac{y}{2}} dy \\ &< 0.9 w^{-\frac{1}{2}} e^{-\frac{w}{2}} \text{ for } w \ge 12,\end{aligned}$$

where the last line is deduced by noticing that the function $h(w) = \int_0^w y^{-\frac{1}{2}} e^{\frac{y}{2}} dy - 2.25 w^{-\frac{1}{2}} e^{\frac{w}{2}}$ satisfies 1) $h(12) < 0$ and 2) $h'(w) = -0.125 w^{-\frac{1}{2}} e^{\frac{w}{2}} + 1.125 w^{-\frac{3}{2}} e^{\frac{w}{2}} = e^{\frac{w}{2}} \cdot w^{-\frac{1}{2}}(1.125 w^{-1} - 0.125) < 0$ for $w \ge 12$. Together they establish $\int_0^w y^{-\frac{1}{2}} e^{\frac{y}{2}} dy < 2.25 w^{-\frac{1}{2}} e^{\frac{w}{2}}$ for $w \ge 12$.

Since there are ξ_1 terms in (i) and ξ_2 terms in (ii), we have

$$\epsilon_U < \xi_1 \big(1 - F_3(w)\big)^2 + 0.9 \xi_2 w^{-\frac{1}{2}} e^{-\frac{w}{2}}. \tag{22}$$

Our next step is to put a lower bound on P_U. According to (17),

$$P_U = P(M) + \sum_{i=1}^{b} P(B_i \cap B_{\mathcal{N}_i}^c \cap M^c). \tag{23}$$

The lower bound of first term in the right side of (23) is obtained by considering only the triplets in $\mathcal{M}$:

$$\begin{aligned}
P(M) &\geq 1 - F_3(w)^{\xi_3} \\
&= 1 - \Big(1 - \big(1 - F_3(w)\big)\Big)^{\xi_3} \\
&\geq \xi_3\big(1 - F_3(w)\big) - \frac{\xi_3(\xi_3 - 1)}{2}\big(1 - F_3(w)\big)^2, \text{ by Taylor expansion} \\
&\geq 0.95\xi_3\big(1 - F_3(w)\big), \text{ as long as } \xi_3\big(1 - F_3(w)\big) < 0.1. \qquad (24)
\end{aligned}$$

Next, we get the lower bound of the second term in the right side of (23). For any fixed i, let Y_1, Y_2 and Y_3 denote the i.i.d. χ_1^2 variables included in the event B_i so that $B_i = \{\sum_{j=1}^3 Y_j > w\}$. Then we have

$$\begin{aligned}
P(B_i \cap B_{\mathcal{N}_i}^c \cap M^c) &= \int_{\sum y_i > w} f_1(y_1)f_1(y_2)f_1(y_3)P(B_{\mathcal{N}_i}^c \cap M^c | y_1, y_2, y_3)d\boldsymbol{y} \\
&= \int_{\sum y_i > w} f_1(\boldsymbol{y})h(\boldsymbol{y}, w)d\boldsymbol{y}
\end{aligned}$$

where we define $f_1(\boldsymbol{y}) = f_1(y_1)f_1(y_2)f_1(y_3)$ and $h(\boldsymbol{y}, w) = P(B_{\mathcal{N}_i}^c \cap M^c | y_1, y_2, y_3)$. In addition, let $V(w) = \{\boldsymbol{y} : \sum_{i=1}^3 y_i > w\}$ denote the region of integration. By Reynolds transport theorem (or multidimensional Leibniz's rule),

$$\begin{aligned}
\frac{d}{dw}P(B_i \cap B_{\mathcal{N}_i}^c \cap M^c) &= \int_{V(w)} f_1(\boldsymbol{y})\frac{\partial}{\partial w}h(\boldsymbol{y}, w)d\boldsymbol{y} \\
&\quad + \int_{\partial V(w)} (\boldsymbol{v}_b \cdot \boldsymbol{n})f_1(\boldsymbol{y})h(\boldsymbol{y}, w)dA \qquad (25)
\end{aligned}$$

where $\partial V(w)$ is the boundary of $V(w)$, $\boldsymbol{v}_b$ is the Eulerian velocity of the boundary, $\boldsymbol{n}$ is the outward-pointing unit-normal, and dA is the surface element.

On the other hand,

$$\begin{aligned}
\frac{d}{dw}P(B_i) &= \frac{d}{dw}\int_{V(w)} f_1(\boldsymbol{y})d\boldsymbol{y} \\
&= \int_{\partial V(w)} (\boldsymbol{v}_b \cdot \boldsymbol{n})f_1(\boldsymbol{y})dA \qquad (26)
\end{aligned}$$

Since both $B_{\mathcal{N}_i}^c$ and M^c strictly enlarges as w increases, $h(\boldsymbol{y}, w)$ is an increasing function of w. This leads to $\frac{\partial}{\partial w}h(\boldsymbol{y}, w) > 0$. Moreover, $h(\boldsymbol{y}, w) < 1$ from definition. Lastly, the fact that $V(w_1) \subset V(w_2)$ for any $w_1 > w_2$ gives $\boldsymbol{v}_b \cdot \boldsymbol{n} \leq 0$. These altogether establish $\frac{d}{dw}P(B_i \cap B_{\mathcal{N}_i}^c \cap M^c) > \frac{d}{dw}P(B_i)$ as we compare the expression in (25) and (26). With the apparent relation $P(B_i \cap B_{\mathcal{N}_i}^c \cap M^c) < P(B_i)$, we conclude that

$$\frac{d}{dw} \cdot \frac{P(B_i \cap B_{\mathcal{N}_i}^c \cap M^c)}{P(B_i)} > 0 \text{ for all } i \qquad (27)$$

Therefore if we only focus on calculating ϵ_U/P_U for $w \geq w_T$ where w_T is a pre-fixed value, then we can first evaluate $\sum_{i=1}^{b} P(B_i \cap B^c_{\mathcal{N}_i} \cap M^c)/\big(1 - F_3(w)\big) = \xi_T$ at $w = w_T$. By (27),

$$\sum_{i=1}^{b} P(B_i \cap B^c_{\mathcal{N}_i} \cap M^c) \geq \xi_T\big(1 - F_3(w)\big) \text{ for } w \geq w_T \tag{28}$$

Plugging in (24) and (28) into (23) gives

$$P_U \geq (0.95\xi_3 + \xi_T)\big(1 - F_3(w)\big) \tag{29}$$

The upper bound on ϵ_U in (22) and the lower bound on P_U in (29) yield

$$\begin{aligned}
\frac{\epsilon_U}{P_U} &< \frac{\xi_1\big(1 - F_3(w)\big)^2 + 0.9\xi_2 w^{-\frac{1}{2}} e^{-\frac{w}{2}}}{(0.95\xi_3 + \xi_T)\big(1 - F_3(w)\big)} \\
&= \frac{\xi_1\big(1 - F_3(w)\big)}{0.95\xi_3 + \xi_T} + \frac{0.9\xi_2 w^{-\frac{1}{2}} e^{-\frac{w}{2}}}{(0.95\xi_3 + \xi_T)\big(1 - F_3(w)\big)} \\
&< \frac{\xi_1\big(1 - F_3(w)\big)}{0.95\xi_3 + \xi_T} + \frac{0.9\xi_2}{0.95\xi_3 + \xi_T}\sqrt{\frac{\pi}{2}} \cdot \frac{1}{w}
\end{aligned} \tag{30}$$

for all $w > w_T \geq 12$. The last line comes from substituting the following equation:

$$1 - F_3(w) = \frac{1}{\sqrt{2\pi}} \int_w^\infty \sqrt{y} e^{-\frac{y}{2}} dy > \frac{\sqrt{w}}{\sqrt{2\pi}} \int_w^\infty e^{-\frac{y}{2}} dy = \sqrt{\frac{2w}{\pi}} e^{-\frac{w}{2}}$$

Lemma 1 in Laurent and Massart (2000) gives $1 - F_3(3 + 2\sqrt{3t} + 2t) \leq e^{-t}$ for all $t > 0$. Since $3 + 2\sqrt{3t} < 4t$ when $t \geq 2$, we have that $1 - F_3(6t) < e^{-t}$ or $1 - F_3(w) < e^{-\frac{w}{6}}$ for $w \geq 12$. Therefore, the rate of decay for ϵ_U/P_U is $\mathcal{O}(e^{-\frac{w}{6}}) + \mathcal{O}(\frac{1}{w})$.

References

Caporaso, J. G., Kuczynski, J., Stombaugh, J., Bittinger, K., Bushman, F. D., Costello, E. K., Fierer, N., Pena, A. G., Goodrich, J. K., Gordon, J. I. et al. (2010). QIIME allows analysis of high-throughput community sequencing data. *Nature Methods* **7** 335-336.

Chen, Y. and Hanson, T. E. (2014). Bayesian nonparametric k-sample tests for censored and uncensored data. *Computational Statistics and Data Analysis* **71** 335-346.

Chen, J. and Li, H. (2013). Variable selection for sparse Dirichlet-multinomial regression with an application to microbiome data analysis. *Annals of Applied Statistics* **7** 418-442.

David, L. A., Maurice, C. F., Carmody, R. C., Gootenberg, D. B., Button, J. E., Wolfe, B. E., Ling, A. V., Devlin, A. S., Varma, Y., Fischbach, M. A. et al. (2014). Diet rapidly and reproducibly alters the human gut microbiome. *Nature* **505** 559-563.

Dohmen, K. (2000). Improved Bonferroni inequalities via union-closed set systems. *Journal of Combinatorial Theory Series A* **92** 61-67.

Dohmen, K. (2003). Improved inclusion-exclusion identities and Bonferroni inequalities with reliability applications. *SIAM Journal on Discrete Mathematics* **16** 156-171.

Dohmen, K. and Tittmann, P. (2004). Bonferroni-Galambos inequalities for partition lattices. *Electronic Journal of Combinatorics* **11** R85.

Efron, B. (1997). The length heuristic for simultaneous hypothesis tests. *Biometrika* **84** 143-157.

Glaz, J., Naus, J. and Wallenstein, S. (2001). Scan statistics. *Springer New York*.

Hahn, T. (2005). Cuba-a library for multidimensional numerical integration. *Computer Physics Communications* **168** 78-95.

Holmes, I., Harris, K. and Quince, C. (2012). Dirichlet multinomial mixtures: generative models for microbial metagenomics. *PLoS ONE* **7** e30126.

Holmes, C. C., Caron, F., Griffin, J. E., and Stephens, D. A. (2015). Two-sample Bayesian nonparametric hypothesis testing. *Bayesian Analysis* **10** 297-320.

Human Microbiome Project Consortium (2012). A framework for human microbiome research. *Nature* **486** 215-221.

Hunter, D. (1976). An upper bound for the probability of a union. *Journal of Applied Probability* **13** 597-603.

Lavine, M. (1992). Some aspects of Polya Tree distributions for statistical modelling. *Annals of Statistics* **20** 1222-1235.

La Rosa, P. S., Brooks, J. P., Deych, E., Boone, E. L., Edwards, D. J., Wang, Q., Sodergren, E., Weinstock, G. and Shannon, W. D. (2012). Hypothesis testing and power calculations for taxonomic-based human microbiome Data. *PLoS ONE* **7** e52078.

Laurent, B. and Massart, P. (2000). Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics* **28** 1302-1338.

McDonald, D., Birmingham, A. and Knight, R. (2015a). Context and the human microbiome. *Microbiome* **3** 1-8.

McDonald, D., Hornig, M., Lozupone, C., Debelius, J., Gilbert, J., and Knight, R. (2015b). Towards large-cohort comparative studies to define the factors influencing the gut microbial community structure of ASD patients. *Microbial Ecology In Health And Disease* **26** 26555.

Mosimann, J. E. (1962). On the compound multinomial distribution, the multivariate β-distribution, and correlations among proportions. *Biometrika* **49** 65-82.

Naiman, D. Q. and Wynn, H. P. (1992). Inclusion-exclusion-Bonferroni identities and inequalities for discrete tube-like problems via Euler characteristics. *Annals of Statistics* **20** 43-76.

Naiman, D. Q. and Wynn, H. P. (1997). Abstract tubes, improved inclusion-exclusion identities and inequalities and importance sampling. *Annals of Statistics* **25** 1954-1983.

Neuman, H,, Debelius, J. W., Knight, R., and Koren, O. (2015). Microbial endocrinology: the interplay between the microbiota and the endocrine system. *FEMS Microbiology Reviews* **39** 509-521.

Soriano, J. and Ma, L. (2016). Probabilistic multi-resolution scanning for two-sample differences *Journal of the Royal Statistical Society: Series B (Statistical Methodology).* doi: 10.1111/rssb.12180 .

Taylor, J. E., Worsley, K. J. and Gosselin, F. (2007). Maxima of discretely sampled random fields, with an application to 'bubbles'. *Biometrika* **94** 1-18.

Turnbaugh, P. J., Ridaura, V. K., Faith, J. J., Rey, F. E., Knight, R. and Gordon, J. I. (2014). The effect of diet on the human gut microbiome: a metagenomic analysis in humanized gnotobiotic Mice. *Science Translational Medicine* **1** 6ra14.

Weir, B. S. and Hill, W. G. (2002). Estimating F-statistics. *Annual Review of Genetics* **36** 721-750.

Ma, L. and Wong, W. H. (2011). Optional Polya tree and Bayesian inference. *Journal Of The American Statistical Association* **106** 15531565.

Worsley, K. J. (1982). An improved Bonferroni inequality and applications. *Biometrika* **69** 297-302.